(11) **EP 1 152 329 A1**(12) **EUROPEAN PATENT APPLICATION**(43) Date of publication:
07.11.2001 Bulletin 2001/45(51) Int Cl.7: **G06F 9/38**(21) Application number: **01302951.7**(22) Date of filing: **29.03.2001**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

- Heintze, Nevin
Morristown, NJ 07960 (US)
- Jeremiasen, Tor E.
Somerset, NJ 08873 (US)
- Kaxiras, Stefanos
Jersey City, NJ 07302 (US)

(30) Priority: **30.03.2000 US 538757**(71) Applicant: **Agere Systems Guardian Corporation**
Orlando, Florida 32819 (US)


(74) Representative: **Williams, David John et al**
Page White & Farrer,
54 Doughty Street
London WC1N 2LS (GB)

(72) Inventors:
• **Berenbaum, Alan David**
New York, NY 10011 (US)

(54) **Method and apparatus for identifying splittable packets in a multithreaded vliw processor**

(57) A method and apparatus are disclosed for allocating functional units in a multithreaded very large instruction word (VLIW) processor. The present invention combines the techniques of conventional very long instruction word (VLIW) architectures and conventional multithreaded architectures to reduce execution time within an individual program, as well as across a workload. The present invention utilizes instruction packet splitting to recover some efficiency lost with conventional multithreaded architectures. Instruction packet splitting allows an instruction bundle to be partially issued in one cycle, with the remainder of the bundle issued during a subsequent cycle. There are times, however, when instruction packets cannot be split without violating the semantics of the instruction packet assembled by the compiler. A packet split identification bit is disclosed that allows hardware to efficiently determine when it is permissible to split an instruction packet. The split bit informs the hardware when splitting is prohibited. The allocation hardware assigns as many instructions from each packet as will fit on the available functional units, rather than allocating all instructions in an instruction packet at one time, provided the split bit has not been set. Those instructions that cannot be allocated to a functional units are retained in a ready-to-run register. On subsequent cycles, instruction packets in which all instructions have been issued to functional units are updated from their thread's instruction stream, while instruction packets with instructions that have been held are retained. The functional unit allocation logic can then

assign instructions from the newly-loaded instruction packets as well as instructions that were not issued from the retained instruction packets.

FIG. 1
PRIOR ART

EP 1 152 329 A1

Description

Cross-Reference to Related Applications

[0001] The present invention is related to United States Patent Application entitled "Method and Apparatus for Allocating Functional Units in a Multithreaded Very Large Instruction Word (VLIW) Processor," Attorney Docket Number (Berenbaum 7-2-3-3); United States Patent Application entitled "Method and Apparatus for Releasing Functional Units in a Multithreaded Very Large Instruction Word (VLIW) Processor," Attorney Docket Number (Berenbaum 8-3-4-4); and United States Patent Application entitled "Method and Apparatus for Splitting Packets in a Multithreaded Very Large Instruction Word (VLIW) Processor," Attorney Docket Number (Berenbaum 9-4-5-5), each filed contemporaneously herewith, assigned to the assignee of the present invention and incorporated by reference herein.

Field of the Invention

[0002] The present invention relates generally to multithreaded processors, and, more particularly, to a method and apparatus for splitting packets in such multithreaded processors.

Background of the Invention

[0003] Computer architecture designs attempt to complete workloads more quickly. A number of architecture designs have been proposed or suggested for exploiting program parallelism. Generally, an architecture that can issue more than one operation at a time is capable of executing a program faster than an architecture that can only issue one operation at a time. Most recent advances in computer architecture have been directed towards methods of issuing more than one operation at a time and thereby speed up the operation of programs. FIG. 1 illustrates a conventional microprocessor architecture 100. Specifically, the microprocessor 100 includes a program counter (PC) 110, a register set 120 and a number of functional units (FUs) 130-N. The redundant functional units (FUs) 130-1 through 130-N provide the illustrative microprocessor architecture 100 with sufficient hardware resources to perform a corresponding number of operations in parallel.

[0004] An architecture that exploits parallelism in a program issues operands to more than one functional unit at a time to speed up the program execution. A number of architectures have been proposed or suggested with a parallel architecture, including superscalar processors, very long instruction word (VLIW) processors and multithreaded processors, each discussed below in conjunction with FIGS. 2, 4 and 5, respectively. Generally, a superscalar processor utilizes hardware at run-time to dynamically determine if a number of operations from a single instruction stream are independent,

and if so, the processor executes the instructions using parallel arithmetic and logic units (ALUs). Two instructions are said to be independent if none of the source operands are dependent on the destination operands of any instruction that precedes them. A very long instruction word (VLIW) processor evaluates the instructions during compilation and groups the operations appropriately, for parallel execution, based on dependency information. A multithreaded processor, on the other hand, executes more than one instruction stream in parallel, rather than attempting to exploit parallelism within a single instruction stream.

[0005] A superscalar processor architecture 200, shown in FIG. 2, has a number of functional units that operate independently, in the event each is provided with valid data. For example, as shown in FIG. 2, the superscalar processor 200 has three functional units embodied as arithmetic and logic units (ALUs) 230-N, each of which can compute a result at the same time.

The superscalar processor 200 includes a front-end section 208 having an instruction fetch block 210, an instruction decode block 215, and an instruction sequencing unit 220 (issue block). The instruction fetch block 210 obtains instructions from an input queue 205 of a single threaded instruction stream. The instruction sequencing unit 220 identifies independent instructions that can be executed simultaneously in the available arithmetic and logic units (ALUs) 230-N, in a known manner. The refine block 250 allows the instructions to complete, and also provides buffering and reordering for writing results back to the register set 240.

[0006] In the program fragment 310 shown in FIG. 3, instructions in locations L1, L2 and L3 are independent, in that none of the source operands in instructions L2 and L3 are dependent on the destination operands of any instruction that precedes them. When the program counter (PC) is set to location L1, the instruction sequencing unit 220 will look ahead in the instruction stream and detect that the instructions at L2 and L3 are independent, and thus all three can be issued simultaneously to the three available functional units 230-N. For a more detailed discussion of superscalar processors, see, for example, James. E. Smith and Gurindar. S. Sohi, "The Microarchitecture of Superscalar Processors," Proc. of the IEEE (Dec. 1995), incorporated by reference herein.

[0007] As previously indicated, a very long instruction word (VLIW) processor 400, shown in FIG. 4, relies on software to detect data parallelism at compile time from a single instruction stream, rather than using hardware to dynamically detect parallelism at run time. A VLIW compiler, when presented with the source code that was used to generate the code fragment 310 in FIG. 3, would detect the instruction independence and construct a single, very long instruction comprised of all three operations. At run time, the issue logic of the processor 400 would issue this wide instruction in one cycle, directing data to all available functional units 430-N. As shown in

FIG. 4, the very long instruction word (VLIW) processor 400 includes an integrated fetch/decode block 420 that obtains the previously grouped instructions 410 from memory. For a more detailed discussion of very long instruction word (VLIW) processors, see, for example, Burton J. Smith, "Architecture and Applications of the HEP Multiprocessor Computer System," SPIE Real Time Signal Processing IV, 241-248 (1981), incorporated by reference herein.

[0008] One variety of VLIW processors, for example, represented by the Multiflow architecture, discussed in Robert P. Colwell et al., "A VLIW Architecture for a Trace Scheduling Compiler," IEEE Transactions on Computers (August 1988), uses a fixed-width instruction, in which predefined fields direct data to all functional units 430-N at once. When all operations specified in the wide instruction are completed, the processor issues a new, multi-operation instruction. Some more recent VLIW processors, such as the C6x processor commercially available from Texas Instruments, of Dallas, TX and the EPIC IA-64 processor commercially available from Intel Corp, of Santa Clara, CA, instead use a variable-length instruction packet, which contains one or more operations bundled together.

[0009] A multithreaded processor 500, shown in FIG. 5, gains performance improvements by executing more than one instruction stream in parallel, rather than attempting to exploit parallelism within a single instruction stream. The multithreaded processor 500 shown in FIG. 5 includes a program counter 510-N, a register set 520-N and a functional unit 530-N, each dedicated to a corresponding instruction stream N. Alternate implementations of the multithreaded processor 500 have utilized a single functional unit 530, with several register sets 520-N and program counters 510-N. Such alternate multithreaded processors 500 are designed in such a way that the processor 500 can switch instruction issue from one program counter/register set 510-N/520-N to another program counter/register set 510-N/520-N in one or two cycles. A long latency instruction, such as a LOAD instruction, can thus be overlapped with shorter operations from other instruction streams. The TERA MTA architecture, commercially available from Tera Computer Company, of Seattle, WA, is an example of this type.

[0010] An extension of the multithreaded architecture 500, referred to as Simultaneous Multithreading, combines the superscalar architecture, discussed above in conjunction with FIG. 2, with the multithreaded designs, discussed above in conjunction with FIG. 5. For a detailed discussion of Simultaneous Multithreading techniques, see, for example, Dean Tullsen et al., "Simultaneous Multithreading: Maximizing On-Chip Parallelism," Proc. of the 22nd Annual Int'l Symposium on Computer Architecture, 392-403 (Santa Margherita Ligure, Italy, June 1995), incorporated by reference herein. Generally, in a Simultaneous Multithreading architecture, there is a pool of functional units, any number of which may

be dynamically assigned to an instruction which can issue from any one of a number of program counter/register set structures. By sharing the functional units among a number of program threads, the Simultaneous Multithreading architecture can make more efficient use of hardware than that shown in FIG. 5.

[0011] While the combined approach of the Simultaneous Multithreading architecture provides improved efficiency over the individual approaches of the superscalar architecture or the multithreaded architecture, Simultaneous Multithreaded architectures still require elaborate issue logic to dynamically examine instruction streams in order to detect potential parallelism. In addition, when an operation takes multiple cycles, the instruction issue logic may stall, because there is no other source of operations available. Conventional multithreaded processors issue instructions from a set of instructions simultaneously, with the functional units designed to accommodate the widest potential issue. A need therefore exists for a multithreaded processor architecture that does not require a dynamic determination of whether or not two instruction streams are independent. A further need exists for a multithreaded architecture that provides simultaneous multithreading. Yet another need exists for a method and apparatus that improve the utilization of processor resources for each cycle.

Summary of the Invention

[0012] Generally, a method and apparatus are disclosed for allocating functional units in a multithreaded very large instruction word (VLIW) processor. The present invention combines the techniques of conventional very long instruction word (VLIW) architectures and conventional multithreaded architectures. The combined architecture of the present invention reduces execution time within an individual program, as well as across a workload. The present invention utilizes instruction packet splitting to recover some efficiency lost with conventional multithreaded architectures. Instruction packet splitting allows an instruction bundle to be partially issued in one cycle, with the remainder of the bundle issued during a subsequent cycle. Thus, the present invention provides greater utilization of hardware resources (such as the functional units) and a lower elapsed time across a workload comprising multiple threads.

[0013] There are times when instruction packets cannot be split without violating the semantics of the instruction packet assembled by the compiler. In particular, the input value of a register is assumed to be the same for instructions in a packet, even if the register is modified by one of the instructions in the packet. If the packet is split, and a source register for one of the instructions in the second part of the packet is modified by one of the instructions in the first part of the packet, then the compiler semantics will be violated.

[0014] Thus, the present invention utilizes a packet split identification bit that allows hardware to efficiently determine when it is permissible to split an instruction packet. Instruction packet splitting increases throughput across all instruction threads, and reduces the number of cycles that the functional units rest idle. The allocation hardware of the present invention assigns as many instructions from each packet as will fit on the available functional units, rather than allocating all instructions in an instruction packet at one time, provided the packet split identification bit has not been set. Those instructions that cannot be allocated to a functional units are retained in a ready-to-run register. On subsequent cycles, instruction packets in which all instructions have been issued to functional units are updated from their thread's instruction stream, while instruction packets with instructions that have been held are retained. The functional unit allocation logic can then assign instructions from the newly-loaded instruction packets as well as instructions that were not issued from the retained instruction packets.

[0015] The present invention utilizes a compiler to detect parallelism in a multithreaded processor architecture. Thus, a multithreaded VLIW architecture is disclosed that exploits program parallelism by issuing multiple instructions, in a similar manner to single threaded VLIW processors, from a single program sequencer, and also supporting multiple program sequencers, as in simultaneous multithreading but with reduced complexity in the issue logic, since a dynamic determination is not required.

[0016] A more complete understanding of the present invention, as well as further features and advantages of the present invention, will be obtained by reference to the following detailed description and drawings.

Brief Description of the Drawings

[0017]

FIG. 1 illustrates a conventional generalized micro-processor architecture;
 FIG. 2 is a schematic block diagram of a conventional superscalar processor architecture;
 FIG. 3 is a program fragment illustrating the independence of operations;
 FIG. 4 is a schematic block diagram of a conventional very long instruction word (VLIW) processor architecture;
 FIG. 5 is a schematic block diagram of a conventional multithreaded processor;
 FIG. 6 illustrates a multithreaded VLIW processor in accordance with the present invention;
 FIG. 7A illustrates a conventional pipeline for a multithreaded processor;
 FIG. 7B illustrates a pipeline for a multithreaded processor in accordance with the present invention;
 FIG. 8 is a schematic block diagram of an imple-

mentation of the allocate stage of FIG. 7B;

FIG. 9 illustrates the execution of three threads TA-TC for a conventional multithreaded implementation, where threads B and C have higher priority than thread A;

FIGS. 10A and 10B illustrate the operation of instruction packet splitting in accordance with the present invention;

FIG. 11 is a program fragment that may not be split in accordance with the present invention;

FIG. 12 is a program fragment that may be split in accordance with the present invention;

FIG. 13 illustrates a packet corresponding to the program fragment of FIG. 12 where the instruction splitting bit has been set;

FIG. 14 is a program fragment that may not be split in accordance with the present invention; and

FIG. 15 illustrates a packet corresponding to the program fragment of FIG. 14 where the instruction splitting bit has not been set.

Detailed Description

[0018] FIG. 6 illustrates a Multithreaded VLIW processor 600 in accordance with the present invention. As shown in FIG. 6, there are three instruction threads, namely, thread A (TA), thread B (TB) and thread C (TC), each operating at instruction number n . In addition, the illustrative Multithreaded VLIW processor 600 includes nine functional units 620-1 through 620-9, which can be allocated independently to any thread TA-TC. Since the number of instructions across the illustrative three threads TA-TC is nine and the illustrative number of available functional units 620 is also nine, then each of the instructions from all three threads TA-TC can issue their instruction packets in one cycle and move onto instruction $n+1$ on the subsequent cycle.

[0019] It is noted that there is generally a one-to-one correspondence between instructions and the operation specified thereby. Thus, such terms are used interchangeably herein. It is further noted that in the situation where an instruction specifies multiple operations, it is assumed that the multithreaded VLIW processor 600 includes one or more multiple-operation functional units 620 to execute the instruction specifying multiple operations. An example of an architecture where instructions specifying multiple operations may be processed is a complex instruction set computer (CISC).

[0020] The present invention allocates instructions to functional units to issue multiple VLIW instructions to multiple functional units in the same cycle. The allocation mechanism of the present invention occupies a pipeline stage just before arguments are dispatched to functional units. Thus, FIG. 7A illustrates a conventional pipeline 700 comprised of a fetch stage 710, where a packet is obtained from memory, a decode stage 720, where the required functional units and registers are identified for the fetched instructions, and an execute

stage 730, where the specified operations are performed and the results are processed.

[0021] Thus, in a conventional VLIW architecture, a packet containing up to K instructions is fetched each cycle (in the fetch stage 710). Up to K instructions are decoded in the decode stage 720 and sent to (up to) K functional units (FUs). The registers corresponding to the instructions are read, the functional units operate on them and the results are written back to registers in the execute stage 730. It is assumed that up to three registers can be read and up to one register can be written per functional unit.

[0022] FIG. 7B illustrates a pipeline 750 in accordance with the present invention, where an allocate stage 780, discussed further below in conjunction with FIG. 8, is added for the implementation of multithreaded VLIW processors. Generally, the allocate stage 780 determines how to group the operations together to maximize efficiency. Thus, the pipeline 750 includes a fetch stage 760, where up to N packets are obtained from memory, a decode stage 770, where the functional units and registers are identified for the fetched instructions (up to $N \times K$ instructions), an allocate stage 780, where the appropriate instructions are selected and assigned to the FUs, and an execute stage 790, where the specified operations are performed and the results are processed.

[0023] In the multithreaded VLIW processor 600 of the present invention, up to N threads are supported in hardware. N thread contexts exist and contain all possible registers of a single thread and all status information required. A multithreaded VLIW processor 600 has M functional units, where M is greater than or equal to K. The modified pipeline stage 750, shown in FIG. 7B, works in the following manner. In each cycle, up to N packets (each containing up to K instructions) are fetched at the fetch stage 760. The decode stage 770 decodes up to $N \times K$ instructions and determines their requirements and the registers read and written. The allocate stage 780 selects M out of (up to) $N \times K$ instructions and forwards them to the M functional units. It is assumed that each functional unit can read up to 3 registers and write one register. In the execute stage 790, up to M functional units read up to $3 \times M$ registers and write up to M registers.

[0024] The allocate stage 780 selects for execution the appropriate M instructions from the (up to) $N \times K$ instructions that were fetched and decoded at stages 760 and 770. The criteria for selection are thread priority or resource availability or both. Under the thread priority criteria, different threads can have different priorities. The allocate stage 780 selects and forwards the packets (or instructions from packets) for execution belonging to the thread with the highest priority according to the priority policy implemented. A multitude of priority policies can be implemented. For example, a priority policy for a multithreaded VLIW processor supporting N contexts (N hardware threads) can have N priority levels. The highest priority thread in the processor is allocated be-

fore any other thread. Among threads with equal priority, the thread that waited the longest for allocation is preferred.

[0025] Under the resource availability criteria, a packet (having up to K instructions) can be allocated only if the resources (functional units) required by the packet are available for the next cycle. Functional units report their availability to the allocate stage 780.

[0026] FIG. 8 illustrates a schematic block diagram of an implementation of the allocate stage 780. As shown in FIG. 8, the hardware needed to implement the allocate stage 780 includes a priority encoder 810 and two crossbar switches 820, 830. Generally, the priority encoder 810 examines the state of the multiple operations in each thread, as well as the state of the available functional units. The priority encoder 810 selects the packets that are going to execute and sets up the first crossbar switch 820 so that the appropriate register contents are transferred to the functional units at the beginning of the next cycle. The output of the priority encoder 810 configures the first crossbar switch 820 to route data from selected threads to the appropriate functional units. This can be accomplished, for example, by sending the register identifiers (that include a thread identifier) to the functional units and letting the functional units read the register contents via a separate data network and using the crossbar switch 810 to move the appropriate register contents to latches that are read by the functional units at the beginning of the next cycle.

[0027] Out of the N packets that are fetched by the fetch stage 760 (FIG. 7B), the priority encoder 810 selects up to N packets for execution according to priority and resource availability. In other words, the priority encoder selects the highest priority threads that do not request unavailable resources for execution. It then sets up the first crossbar switch 810. The input crossbar switch 810 routes up to $3K \times N$ inputs to up to $3 \times M$ outputs. The first crossbar switch 810 has the ability to transfer the register identifiers (or the contents of the appropriate registers) of each packet to the appropriate functional units.

[0028] Since there are up to N threads that can be selected in the same cycle and each thread can issue a packet of up to K instructions and each instruction can read up to 3 registers there are $3K \times N$ register identifiers to select from. Since there are only M functional units and each functional unit can accept a single instruction, there are only 3M register identifiers to be selected. Therefore, the crossbar switch implements a $3K \times N$ to 3M routing of register identifiers (or register contents).

[0029] The output crossbar switch 830 routes M inputs to $N \times M$ or $N \times K$ outputs. The second crossbar switch 830 is set up at the appropriate time to transfer the results of the functional units back to the appropriate registers. The second crossbar switch 830 can be implemented as a separate network by sending the register identifiers (that contain a thread identifier) to the functional units. When a functional unit computes a result,

the functional unit routes the result to the given register identifier. There are M results that have to be routed to up to N threads. Each thread can accept up to K results. The second crossbar switch 830 routes M results to N*K possible destinations. The second crossbar switch 830 can be implemented as M buses that are connected to all N register files. In this case, the routing becomes M results to N*M possible destinations (if the register files have the ability to accept M results).

[0030] In a conventional single-threaded VLIW architecture, all operations in an instruction packet are issued simultaneously. There are always enough functional units available to issue a packet. When an operation takes multiple cycles, the instruction issue logic may stall, because there is no other source of operations available. In a multithreaded VLIW processor in accordance with the present invention, on the other hand, these restrictions do not apply.

[0031] FIG. 9 illustrates the execution of three threads TA-TC for a conventional multithreaded implementation (without the benefit of the present invention), where threads B and C have higher priority than thread A. Since thread A runs at the lowest priority, its operations will be the last to be assigned. As shown in FIG. 9, five functional units 920 are assigned to implement the five operations in the current cycle of the higher priority threads TB and TC. Thread A has four operations, but there are only two functional units 920 available. Thus, thread A stalls for a conventional multithreaded implementation.

[0032] In order to maximize throughput across all threads, and minimize the number of cycles that the functional units rest idle, the present invention utilizes instruction packet splitting. Instead of allocating all operations in an instruction packet at one time, the allocation hardware 780, discussed above in conjunction with FIG. 8, assigns as many operations from each packet as will fit on the available functional units. Those operations that will not fit are retained in a ready-to-run register 850 (FIG. 8). On subsequent cycles, instruction packets in which all operations have been issued to functional units are updated from their thread's instruction stream, while instruction packets with operations that have been held are retained. The functional unit allocation logic 780 can then assign operations from the newly-loaded instruction packets as well as operations that were not issued from the retained instruction packets.

[0033] The operation of instruction packet splitting in accordance with the present invention is illustrated in FIGS. 10A and 10B. In FIG. 10A, there are three threads, each with an instruction packet from location *n* ready to run at the start of cycle *x*. Thread A runs at the lowest priority, so its operations will be the last to be assigned. Threads B and C require five of the seven available functional units 1020 to execute. Only two functional units 1020-2 and 1020-6 remain, so only the first two operations from thread A are assigned to execute. All

seven functional units 1020 are now fully allocated.

[0034] At the completion of cycle *x*, the instruction packets for threads B and C are retired. The instruction issue logic associated with the threads replaces the instruction packets with those for address *n+1*, as shown in FIG. 10B. Since the instruction packet for thread A is not completed, the packet from address *n* is retained, with the first two operations marked completed. On the next cycle, *x+1*, illustrated in FIG. 10B, the final two operations from thread A are allocated to functional units, as well as all the operations from threads B and C. Thus, the present invention provides greater utilization of hardware resources (i.e., the functional units 1020) and a lower elapsed time across a workload comprising the multiple threads.

[0035] An instruction packet cannot be split without verifying that splitting would violate the semantics of the instruction packet assembled by the compiler. In particular, the input value of a register is assumed to be the same for instructions in a packet, even if the register is modified by one of the instructions in the packet. If the packet is split, and a source register for one of the instructions in the second part of the packet is modified by one of the instructions in the first part of the packet, then the compiler semantics will be violated. This is illustrated in the program fragment 1110 in FIG. 11:

[0036] As shown in FIG. 11, if instructions in locations L1, L2 and L3 are assembled into an instruction packet, and R0 = 0, R1 = 2, and R2 = 3 before the packet executes, then the value of R0 will be 5 after the packet completes. On the other hand, if the packet is split and instruction L1 executes before L3, then the value of R0 will be 2 after the packet completes, violating the assumptions of the compiler.

[0037] One means to avoid packet splitting that would violate program semantics is to add hardware to the instruction issue logic that would identify when destination registers are used as sources in other instructions in an instruction packet. This hardware would inhibit packet splitting when one of these read-after-write hazards exists. The mechanism has the disadvantage of requiring additional hardware, which takes area resources and could possibly impact critical paths of a processor and therefore reduce the speed of the processor.

[0038] A compiler can easily detect read-after-write hazards in an instruction packet. It can therefore choose to never assemble instructions with these hazards into an instruction packet. The hardware would then be forced to run these instructions serially, and thereby avoid the hazard. Any instruction packet that had a read-after-write hazard would be considered in error, and the architecture would not guarantee the results. Although safe from semantic violations, this technique has the disadvantage that it does not exploit potential parallelism in a program, since the parallelism available in the instruction packet with a hazard is lost, even if the underlying hardware would not have split the packet.

[0039] The present invention combines compiler

knowledge with a small amount of hardware. In the illustrative implementation, a single bit, referred to as the split bit, is placed in the prefix of a multi-instruction packet to inform the hardware that this packet cannot be split. Since the compiler knows which packets have potential read-after-write hazards, it can set this bit in the packet prefix whenever a hazard occurs. At run-time, the hardware will not split a packet with the bit set, but instead will wait until all instructions in the packet can be run in parallel. This concept is illustrated in FIGS. 12 through 15.

[0040] The compiler detects that the three instruction sequence 1210 in FIG. 12 can be split safely, so the split bit is set to 1, as shown in FIG. 13. In FIG. 14, on the other hand, the three instruction sequence 1410 cannot be split, since there is a read-after-write hazard between instructions L1 and L3. The split bit is therefore set to 0, as shown in FIG. 15

[0041] It is to be understood that the embodiments and variations shown and described herein are merely illustrative of the principles of this invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention.

Claims

1. A multithreaded very large instruction word (VLIW) processor, comprising:
 - a plurality of functional units for executing a plurality of instructions from a multithreaded instruction stream, said instructions being grouped into packets by a compiler, said compiler including an indication in said packet of whether said instructions in said packet may be split; and
 - an allocator that selects instructions from said instruction stream and forwards said instructions to said plurality of functional units, said allocator assigning instructions from at least one of said instruction packets to a plurality of said functional units if said indication indicates said packet may be split.
2. The multithreaded very large instruction word (VLIW) processor of claim 1, wherein said indication is a split bit.
3. The multithreaded very large instruction word (VLIW) processor of claim 1, wherein said allocator assigns as many instructions from a given instruction packet as permitted by an availability of said functional units.
4. The multithreaded very large instruction word (VLIW) processor of claim 1, further comprising a

register for storing for execution in a later cycle an indication of those instructions from a given instruction packet that cannot be allocated to a functional unit in a given cycle.

5. The multithreaded very large instruction word (VLIW) processor of claim 4, wherein instruction packets in which all instructions have been issued to functional units are updated from the instruction stream of said thread.
6. The multithreaded very large instruction word (VLIW) processor of claim 4, wherein instruction packets with instructions indicated in said register are retained.
7. A method of processing instructions from a multithreaded instruction stream in a multithreaded very large instruction word (VLIW) processor, comprising the steps of:
 - executing said instructions using a plurality of functional units, said instructions being grouped into packets by a compiler, said compiler including an indication in said packet of whether said instructions in said packet may be split; and
 - assigning instructions from at least one of said instruction packets to a plurality of said functional units if said indication indicates said packet may be split; and
 - forwarding said selected instructions to said plurality of functional units.
8. The method of claim 7, wherein said indication is a split bit.
9. The method of claim 7, wherein said assigning step assigns as many instructions from a given instruction packet as permitted by an availability of said functional units.
10. The method of claim 7, further comprising the step of storing for execution in a later cycle an indication of those instructions from a given instruction packet that cannot be allocated to a functional unit in a given cycle.
11. The method of claim 10, wherein instruction packets in which all instructions have been issued to functional units are updated from the instruction stream of said thread.
12. The method of claim 10, wherein instruction packets with instructions indicated in said register are retained.
13. An article of manufacture for processing instruc-

tions from an instruction stream having a plurality of threads in a multithreaded very large instruction word (VLIW) processor, comprising:

a computer readable medium having computer readable program code means embodied thereon, said computer readable program code means comprising program code means for causing a computer to:

execute said instructions using a plurality of functional units, said instructions being grouped into packets by a compiler, said compiler including an indication in said packet of whether said instructions in said packet may be split; and
assign instructions from at least one of said instruction packets to a plurality of said functional units if said indication indicates said packet may be split; and
forward said selected instructions to said plurality of functional units.

14. A compiler for a multithreaded very large instruction word (VLIW) processor, comprising:

a memory for storing computer-readable code; and
a processor operatively coupled to said memory, said processor configured to:

translate instructions from a program into a machine language;
group a plurality of said instructions into a packet; and
provide an indication with said packet indicating whether said instructions in said packet may be split.

15. The compiler of claim 14, wherein said instruction packet can be split provided the semantics of the instruction packet assembled by the compiler are not violated.

16. The compiler of claim 14, wherein said instruction packet can be split provided a source register for one of the instructions in a first part of said packet is not modified by one of the instructions in a second part of said packet.

17. An article of manufacture for compiling instructions from an instruction stream having a plurality of threads for use in a multithreaded very large instruction word (VLIW) processor, comprising:

a computer readable medium having computer readable program code means embodied thereon, said computer readable program code means comprising program code means for causing a computer to:

translate instructions from a program into a machine language;
group a plurality of said instructions into a packet; and
provide an indication with said packet indicating whether said instructions in said packet may be split.

FIG. 1
PRIOR ART




FIG. 2
PRIOR ART




FIG. 3

310 { L3: ADD R0, R1, R2
L2: SUB R3, R4, R2
L1: OR R6, R1, R5

FIG. 4
PRIOR ART

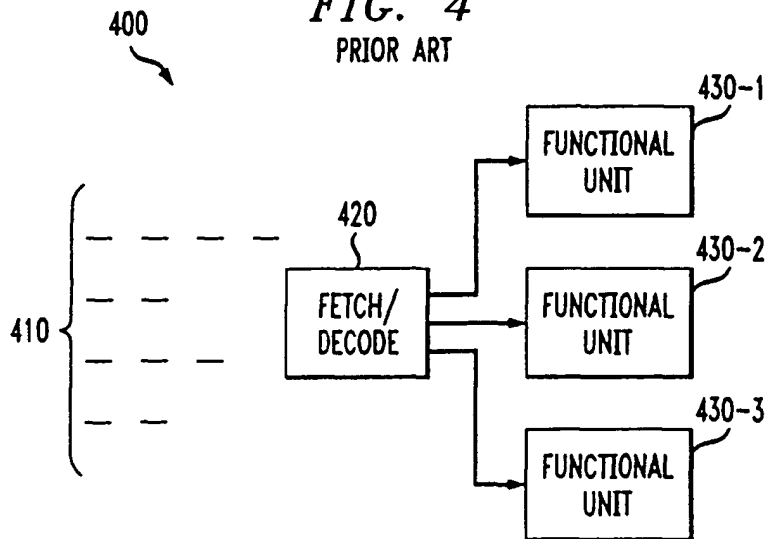


FIG. 5
PRIOR ART




FIG. 6




FIG. 7A

PRIOR ART

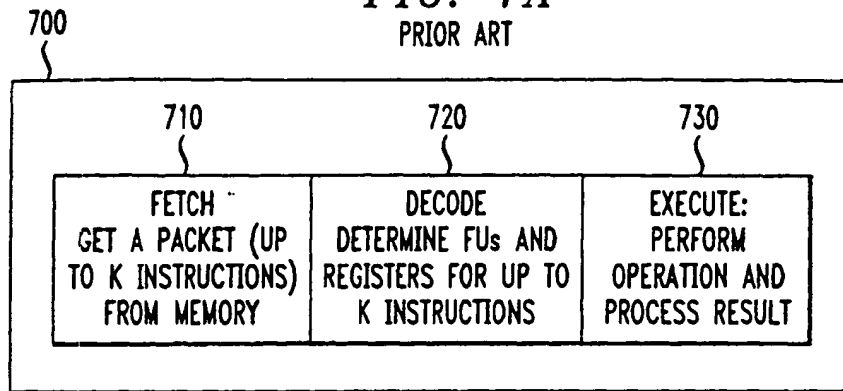
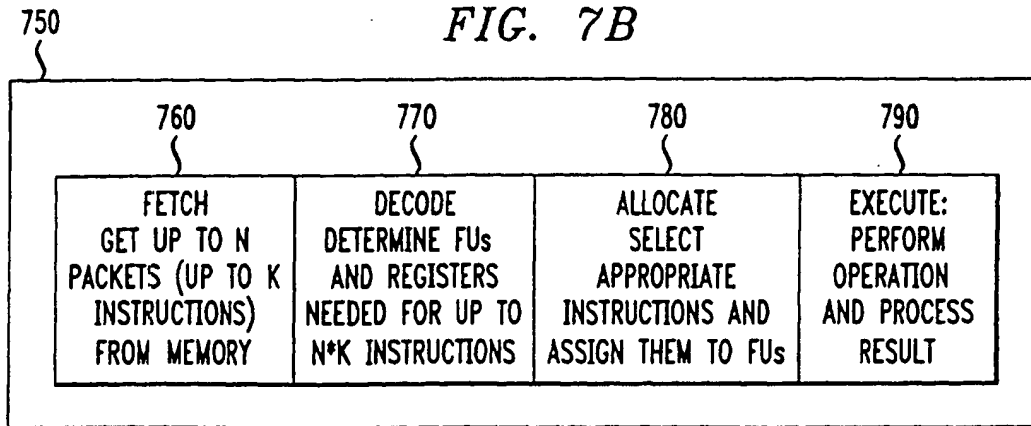


FIG. 7B



780

FIG. 8




FIG. 9

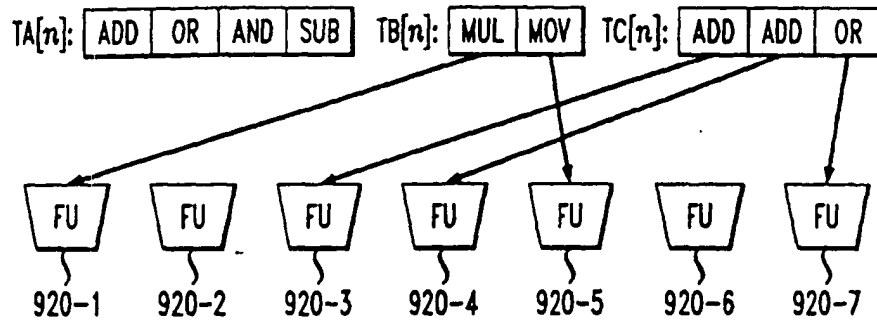


FIG. 10A

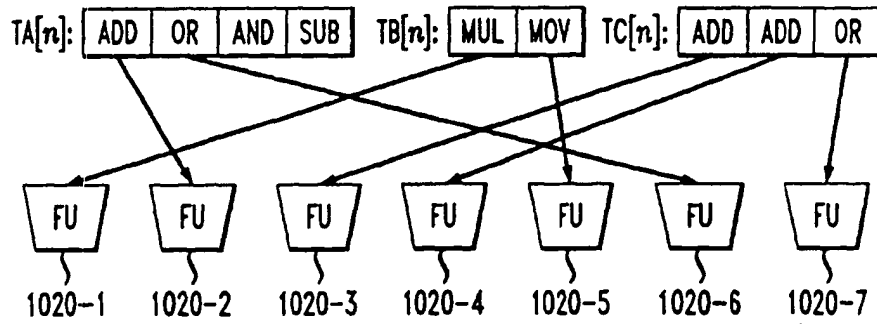


FIG. 10B

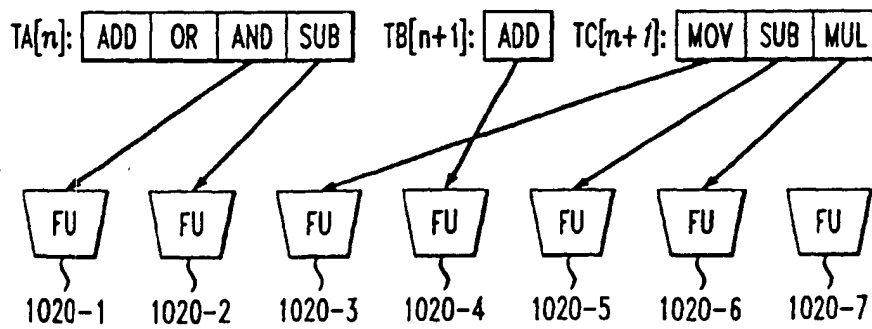


FIG. 11

1110 { L3: ADD R0, R1, R2
L2: SUB R3, R4, R2
L1: MOV R1, R0

FIG. 12

1210 { L3: ADD R6, R1, R2
L2: SUB R3, R4, R2
L1: MOV R5, R0

FIG. 13





FIG. 14

1410 { L3: ADD R0, R1, R2
L2: SUB R3, R4, R2
L1: MOV R1, R0

FIG. 15





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 2951

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.CI.7)
A	US 5 574 939 A (KECKLER STEPHEN W ET AL) 12 November 1996 (1996-11-12) * abstract * * column 3, line 58 - column 4, line 8; figure 1 * * column 8, line 25 - line 57 * * column 9, line 29 - column 10, line 17 *	1-17	606F9/38
A	US 5 404 469 A (CHUNG JIN-CHIN ET AL) 4 April 1995 (1995-04-04) * abstract * * column 6, line 35 - column 8, line 55 *	1-17	
P,A	EP 1 050 808 A (ST MICROELECTRONICS SA) 8 November 2000 (2000-11-08) * paragraph '0006! - paragraph '0009! * * paragraph '0032! - paragraph '0035! *	1,7,13, 14,17	
P,A	US 6 170 051 B1 (DOWLING ERIC M) 2 January 2001 (2001-01-02) * column 5, line 11 - line 38 * * column 10, line 56 - column 11, line 20 * * column 13, line 11 - line 30 *	1-17	
			TECHNICAL FIELDS SEARCHED (Int.CI.7)
			606F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 4 September 2001	Examiner Moraiti, M
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPF FORM 1503 03 82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 30 2951

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

04-09-2001

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US 5574939	A	12-11-1996	WO	9427216 A	24-11-1994
US 5404469	A	04-04-1995	DE	4217012 A	26-08-1993
			JP	2928695 B	03-08-1999
			JP	7191847 A	28-07-1995
EP 1050808	A	08-11-2000	JP	2000330790 A	30-11-2000
US 6170051	B	02-01-2001	US	6163836 A	19-12-2000

EPC FORM P4439

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82